4DVAR Assimilation of ADCP Data with the Navy Coastal Ocean Model using the Cycling Representer Method

S. R. Smith, H. E. Ngodock, and G. A. Jacobs
Naval Research Laboratory
Code 7320, Bldg. 1009
Stennis Space Center, MS 39529 USA

ABSTRACT - **4D-variational assimilation (4DVAR) is used to combine ADCP velocity observations with the Navy Coastal Ocean model (NCOM) to obtain an optimal solution that minimizes a cost function containing the weighted squared errors of velocity measurements, initial conditions, boundary conditions, and model dynamics. However, in order to converge to the global minimum of this cost function, the ocean model (and its adjoint) must be linear. Ocean models, especially those that are designed to resolve baroclinic and mesoscale processes, are typically highly-nonlinear and must be linearized. Tangent linearization is a linearization method that is performed by expanding the nonlinear dynamics about a background field using the first order approximation of Taylor's expansion. The accuracy and stability of this tangent linearized model (TLM) is a very sensitive function of the background accuracy, the level of nonlinearity of the model, complexity of the bathymetry, and the complexity of the flow field. Therefore, in high-resolution coastal domains, the TLM is only going to be stable for a relatively short period of time.**

In this paper, assimilation experiments are performed in a high-resolution Mississippi Bight coastal domain. The TLM of NCOM for this domain is only accurate for about 1 day. The representer method is used to solve this highly nonlinear, weak-constraint, 4DVAR problem. However, due to the short stability time period of this assimilation problem, the representer method is cycled by splitting the time period of the assimilation problem into smaller cycles, therefore ensuring TLM stability and proper data assimilation. The cycle time period needs to be such that it is short enough for the TLM to be stable, but long enough to minimize the loss of information due to reducing the temporal correlation of the dynamic error. We have found that for the Mississippi Bight experiments presented in this paper that a cycle length of 1 day works best. For each new cycle, a background is first created as a nonlinear forecast from the previous cycle's assimilated solution. Then, data that falls within the time period of the new cycle is used to calculate a new assimilated solution using the previous cycle's forecast as the background.

The cycling representer method has been previously demonstrated to drastically improve assimilation accuracy with simpler nonlinear models. Now, this assimilation method is being applied to NCOM. This assimilation system is demonstrated in the Mississippi Bight by assimilating velocity measurements from an array of 14 ADCP moorings for the month of June, 2004. The initial condition for the first cycle, the boundary conditions, and the background around which the TLM is expanded comes from an operational global NCOM. The weak-constraint cycling representer method corrects the velocity components of initial conditions, boundary conditions, and dynamics. This paper will demonstrate the improvement of assimilation accuracy as the time window of the cycles is reduced to 1 day, but when 12-hour cycles are used, the system begins to lose skill. It will also be demonstrated that the forecast skill will be improved as the assimilation system progresses through the cycles.

I. INTRODUTION

Improving the capability to model and forecast the fundamental properties of the ocean has been the endeavor of numerous universities, institutions, and agencies for many decades. These institutions have their reasons for wanting the increased modeling capability and their reasons require resolution with very large spatial and temporal scales (global) all the way down to small scales (littoral, coastal regions). The primary oceanic processes that govern the motion and physics of the ocean are going to be drastically different between these extreme differences in scales. For example, the ability to model and predict the ocean in a small scale coastal domain is going to require a different strategy than say a large scale open ocean domain. Processes governing the oceanic properties in the small-scale coastal domain will rely more heavily on the bathymetry, coastal geometry, mixing, river inflow, and other nonlinear (NL) interactions. Therefore, it is safe to say that the system required to accurately model and forecast in such a region will generally need to be more sophisticated.

One of the key components to accurately forecasting the ocean, regardless of the resolution scale, is data, and how best to merge the data with the model. Since models are always going to be in error and data is almost always going to be scarce relative to a discretized model, the big question is how does one extract the most information possible from the data and apply it to maximize its influence on the model. Simpler data assimilation techniques such as nudging, optimal interpolation, and even 3DVAR have a limited range of influence on the model state (especially in the temporal dimension) and they do not take into account the physics of the model. These techniques may be adequate and desirable for large-scale open ocean models since they are computationally cheaper relative to the more sophisticated techniques, and the open ocean is generally less dynamic and more linear, therefore causing the model to be more accurate in this region. Ocean dynamics, however, become increasingly complex in coastal and littoral regions due to the larger influence of baroclinic and NL processes as well as the complex geometry and interactions with land. Computationally intensive assimilation techniques may not be feasible for very large domains, but in small

0-933957-35-1

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
30-01-2008	Conference Proceeding	

4. TITLE AND SUBTITLE
4DVAR Assimilation of ADCP Data with the Navy Coastal Ocean Model using the Cycling Representer Method

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
0601153N

6. AUTHOR(S)
Scott R. Smith, Hans E. Ngodock, Gregg A. Jacobs

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER
73-8554-07-5

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Naval Research Laboratory
Oceanography Division
Stennis Space Center, MS 39529-5004

8. PERFORMING ORGANIZATION REPORT NUMBER
NRL/PP/7320-07-7248

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Office of Naval Research
800 N. Quincy St.
Arlington, VA 22217-5660

10. SPONSOR/MONITOR'S ACRONYM(S)
ONR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
4D-variational assimilation (4DVAR) is used to combine ADCP velocity observations with the Navy Coastal Ocean model (NCOM) to obtain an optimal solution that minimizes a cost function containing the weighted squared errors of velocity measurements, initial conditions, boundary conditions, and model dynamics. However, in order to converge to the global minimum of this cost function,the ocean model (and its adjoint) must be linear. Ocean models, especially those that are designed to resolve baroclinic and mesoscale processes, are typically highly-nonlinear and must be linearized. Tangent linearization is a linearization method that is performed by expanding the nonlinear dynamics about a background field using the first order approximation of Taylor's expansion. The accuracy and stability of this tangent linearized model (TLM) is a very sensitive function of the background accuracy, the level of nonlinearity of the model, complexity of the bathymetry, and the complexity of the flow field. Therefore, in high-resolution coastal domains, the TLM is only going to be stable for a relatively short period of time.
In this paper, assimilation experiments are performed in a high-resolution Mississippi Bight coastal domain. The TLM of NCOM for this domain is only accurate for about 1 day. The representer method is used to solve this highly nonlinear, weak-constraint, 4DVAR problem. However, due to the short..

15. SUBJECT TERMS
tangent linearized model, 4DVAR, NCOM

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Scott R. Smith
Unclassified	Unclassified	Unclassified	UL	13	19b. TELEPHONE NUMBER *(Include area code)* 228-688-4630

coastal regions, advanced data assimilation is not only feasible, it is crucial in order to properly account for the complex physics in these types of regions. Advanced sequential data assimilation techniques such as the various Kalman filters take into account the model physics and can propagate the influence of the data forward in time through the model. However, with the rapid increase in computational resources, 4-dimensional variational assimilation (4DVAR) is rapidly becoming a popular assimilation technique since it not only can propagate the influence of data forward in time but also backwards in time (through the adjoint) and can correct the initial conditions of the model.

One of the more sophisticated tools that can be used to solve 4DVAR problems is the representer method developed by Andrew Bennett from Oregon State University ([2], [3], [4], & [5]), and is currently being employed by several institutions whom are developing assimilation systems with complex, baroclinic ocean models. For example, [8] and [9] are using the representer method to assimilate data with ROMS and ADCIRC respectively. One of the biggest challenges in developing a representer based assimilation system is that the model and its adjoint must first be linearized about a particular background. Since the linearized version of the model (TLM) is a simplification of the full NL version, it will only be accurate and numerically stable for a finite period of time. There are many factors that contribute to how long the TLM can remain stable, such as the level of nonlinearity of the model, the accuracy of the background, and the complexity of the bathymetry. Needless to say that for a small-scale coastal application (which is an ideal application for this assimilation method), the time period of TLM stability is going to be relatively short.

The cycling representer method was first introduced and applied by [16], and can be used to overcome the problem of TLM stability. Ref. [16] cycled the representer method in time with a linear 1-dimensional transport model in order to lay down the concept. Then in [17] the authors applied the cycling representer method to a linear, barotropically, unstable shallow water system. In these two applications there was no issue with the TLM since both models were linear. More recently, [12] and [13] explored the idea of applying the cycling representer method to the NL Lorenz attractor and reduced gravity ocean models respectively. In these two papers, an initial background is first created by propagating the NL model forward over the first cycle time period. Then the TLM and adjoint of these perspective models is used to perform an assimilation with the representer method. For the second cycle, a background (forecast) is created by using the final assimilated solution from the first cycle and propagating the NL model forward over the second cycle time period, and then performing an assimilation using this new background. This process is repeated for all subsequent cycles. Ref. [12] and [13] demonstrate that the cycling representer method can be extremely beneficial in situations where the TLM is not stable for a long enough period of time. It is shown that cycling not only eliminates the difficulties associated with an unstable TLM, it significantly reduces the overall cost of the assimilation, particularly when the need for outer loops is dropped. Outer loops are typically required for NL models; they are iterations over the linearizations of the NL Euler-Lagrange problem associated with the minimization of the cost function involving the NL model [9]. By using the Cycling Representer Method, there is a good chance that the need for outer loops can be dropped, because the background is being updated in each cycle. Ref. [12] and [13] demonstrate that once the system is spun-up (typically after the first few cycles), the background is trained towards the data and the TLM is accurate enough to eliminate the need of outer iterations.

Previous applications of the cycling representer method have been with either linear models or simple, low-dimensional NL problems. In this paper, the validity of the cycling representer method is taken one step forward and demonstrated in a realistic application with the Navy Coastal Ocean Model (NCOM). NCOM is a NL, multi-layered, baroclinic ocean model designed to resolve coastal features [1]. Experiments are performed by assimilating velocity measurements from an array of 14 ADCP moorings deployed on the shelf and slope of the Mississippi Bight during the month of June 2004 (Fig. 1). The data set and region used in these experiments are optimal in demonstrating the importance and uniqueness that this assimilation technique offers. This is because velocity measurements have always been notoriously difficult to assimilate into ocean models [14]. This is especially the case in highly dynamical shelf-break regions such as the Mississippi Bight where the circulation is dominated by multiple processes such as inertial oscillations, winds, and intruding eddies [15]. All ocean models (including NCOM) have a difficult time accurately accounting for all of the dynamics in a shelf-break region and matching a high-resolution velocity data set (such as the one used in this study). Prior to the assimilation experiments, NCOM results were compared to the data set, and there were many features observed in the data that were not resolved by the model. Most assimilation techniques can not handle this discrepancy, and would produce dynamically inconsistent results. The cycling representer method is unique in that by assimilating over shorter periods of time and continuously updating the background, there is an improved capability of keeping the solution inline with the data and increasing its accuracy with subsequent cycles. More importantly though is that initial condition, boundary condition, and dynamic error covariances can be specified to account for the unresolved model dynamics. Therefore, during the cost-function minimization, the dynamics, along with the initial and boundary conditions, will be modified over the entire space and time domain to produce a dynamically consistent solution (within the specified error limits) that best matches the data.

This paper is a continuation of previous work ([12] and [13]) and is an ongoing effort to demonstrate the importance and usefulness of applying the cycling representer method to an operational assimilation system for coastal regions that can assimilate a wide variety of measurement types and produce accurate forecasts in real-time. In the next section of this paper, the setup of the experiment will be described. This includes a description of the domain, the model, the data, and the assimilation technique. In

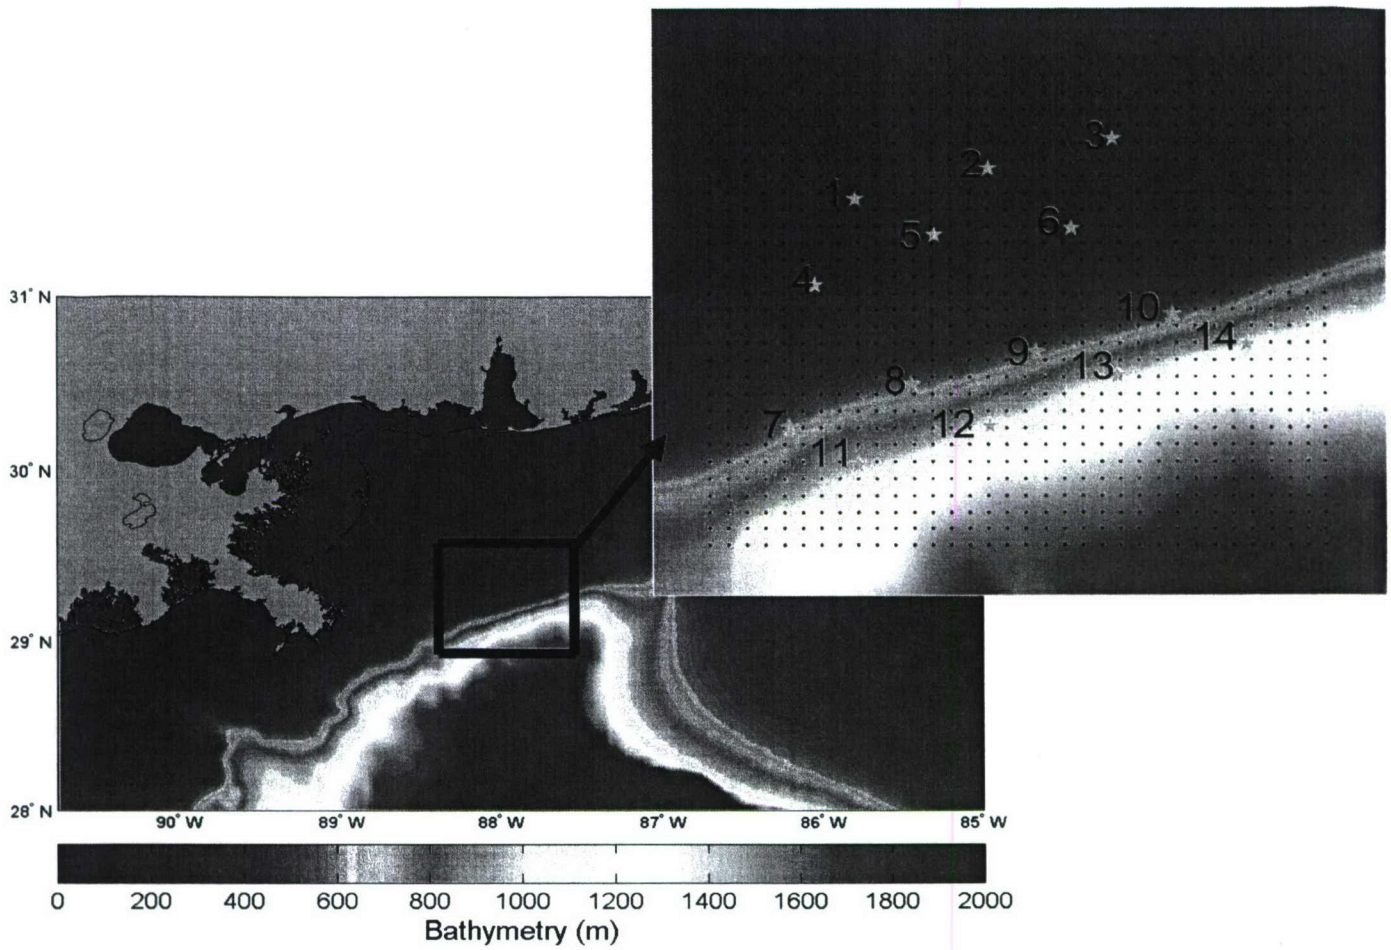

Fig. 1 The Mississippi Bight domain used for this study (black box). The 30X34 black dots represent the discretized grid of the model and the 14 numbered gray stars represent the ADCP mooring locations.

the third section, the assimilation results will be presented and discussed. Then, the final section will contain some concluding remarks.

II. Experiment Setup

A. The Nonlinear Forward Ocean Model

The forward ocean model used in this assimilation experiment is NCOM. NCOM is a free-surface model based on the primitive equations and the hydrostatic, Boussinesq, and incompressible approximations. In many ways NCOM is similar to the Princeton Ocean Model (POM), in that it uses an Arakawa C-grid, is leapfrog in time with an Asselin filter, and uses Smagorinsky coefficients and the Mellor-Yamada 2.5 turbulent closure scheme to parameterize the horizontal and vertical mixing coefficients respectively. NCOM is different from POM primarily in that the free-surface is computed implicitly and NCOM uses both sigma coordinates for the upper layers and z-level coordinates for the lower layers. Further details of NCOM can be found in [1]. Fig. 1 displays the NCOM grid that will be used for the NL forward model, the TLM, and the adjoint model needed for the assimilation experiments in this paper. The 30X34 black dots represent the center points of the Arakawa C-grid and are spaced 2.5km apart, requiring a 4 minute time-step for numerical stability. In the vertical, there are 40 layers with 19 sigma layers in the upper 137m to resolve the shelf-break.

The forward ocean model has several distinct purposes within the cycling representer system. First, a global solution is needed for the initial conditions for the first cycle, and the open boundary conditions for each time-step of the entire time period. In regards to this paper, the phrase global solution means a solution that is larger than the assimilation domain. For these experiments, historical results are extracted from the operational 1/8° global NCOM [1] and used for this purpose. Since the horizontal and temporal resolution of the historical global NCOM solutions are significantly lower than what is needed for these experiments, they are linearly interpolated to the experiment domain and time-steps. Vertical interpolation is not needed because the Mississippi Bight domain uses the same vertical structure. The second purpose of NCOM is to use the initial and boundary conditions from the above global solutions and propagate the NL model forward over each of the cycle time periods

for the Mississippi Bight domain to create a background solution for each cycle. Note that the global initial conditions are only needed for the first cycle; subsequent cycles use the final assimilated solution from the previous cycle as the initial conditions for the NL background forecast.

B. The Tangent Linearized Ocean Model and its adjoint

Since the Representer assimilation technique is essentially performed by minimizing a quadratic cost-function, the model must be first linearized in order for an absolute minimum of the cost function to be determined. This is accomplished by using the first-order approximation of Taylor's expansion of the NCOM dynamics expanded about the above mentioned background.

$$\mathbf{A}_{TL}\vec{x} = \mathbf{A}\vec{x}_{BG} + \frac{d\left(\mathbf{A}\vec{x}_{BG}\right)}{d\vec{x}_{BG}}\left(\vec{x} - \vec{x}_{BG}\right) \tag{1}$$

In the above equation, **A** is the NL model, $\vec{x}$ is the solution state, and $\vec{x}_{BG}$ is the background state. The accuracy of the TLM is plotted for the first 10 days in Fig. 2. TLM accuracy is defined as the root-mean-square (RMS) of the velocity difference between the NL background and tangent linearized solutions. As can be seen, the TLM is fairly accurate for only about a day, which is a relatively short period of time. There are many reasons why the TLM stability is short. The first and foremost reason is that NCOM contains many highly non-linear dynamical components that require linearization such as: the Smagorinsky horizontal mixing scheme, the Coriolis and curvature terms, density, advection, and every time depth is used (in NCOM, depth includes SSH). In these experiments, bottom friction and the Mellor-Yamada vertical mixing scheme are not linearized and are based fully on the background state. The inclusion of these two linearized components was attempted, but the TLM was too unstable in order to perform a reasonable assimilation experiment. The second important criterion that impacts the TLM stability is the accuracy of the background solution. Global NCOM has roughly 1/5 the horizontal resolution and 1/90 the temporal resolution (global NCOM solutions are archived every 6hr) relative to the local Mississippi Bight model. Therefore, since the background is derived from the initial and boundary conditions of Global NCOM, the initial background is going to be in significant error; it is smooth and lacks small-scale features. Another important condition that is reducing the TLM stability is the steep bathymetry in the southeast corner of the domain. This steep bathymetry is amplifying the NL baroclinic processes occurring at the shelf-break.

Another key component of the Representer method is of course the adjoint. If the TLM is expressed in matrix form, as in (1), then the adjoint is simply the transpose of the TLM, $\left(\mathbf{A}_{TL}\vec{x}\right)^{T}$. The TLM of NCOM, however, is not in matrix form and the adjoint was manually generated by reversing all of the operations in space and time.

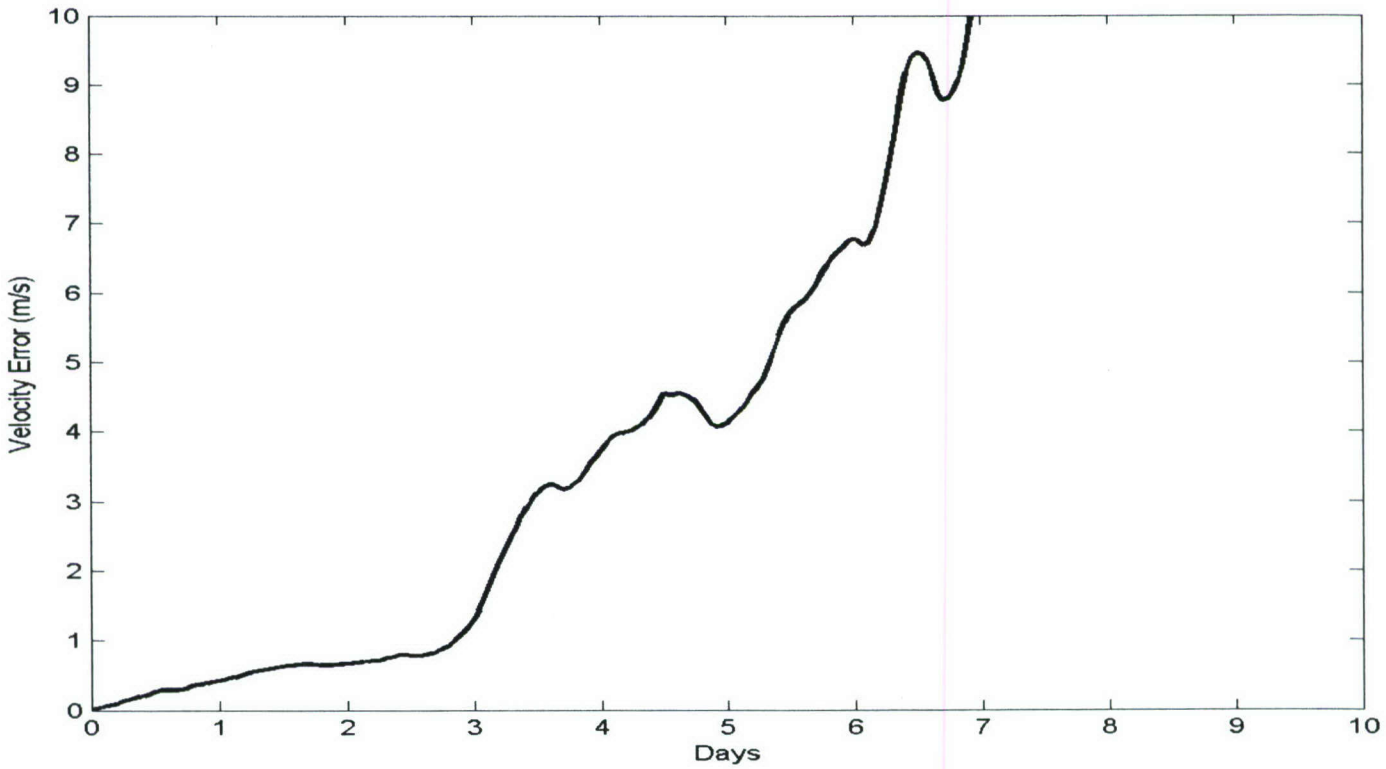

Fig 2 TLM stability is represented here by the evolution of the RMS of the velocity difference between the tangent linearized and nonlinear solutions of NCOM. The TLM maintains relative accuracy for only about the first day, and then the error of the TLM surpasses 1 m/s after about 3 days of propagation.

C. ADCP Data

In May 2004, the Naval Research Laboratory (NRL) deployed an array of 14 ADCP moorings for 1 year along the shelf, shelf-break, and slope of the Mississippi Bight (about 100 miles south of Mobile, Alabama). These moorings were spaced about 10-20 km apart and are identified in Fig. 1 as numbered grey stars. Moorings 1-6 consisted of trawl resistant bottom mounted ADCPs located along the 60m and 90m isobaths. The measurements collected from these shelf moorings were binned into 15 minute time intervals and 2m depth intervals. The slope moorings (7-14) consisted of long-range ADCPs buoyed at 500m depth and located along the 500m and 1000m isobaths. Velocity measurements for the upper 500m were collected and binned in 1 hour time intervals and 10m depth intervals. Even though the tides in this region have relatively small amplitudes, the tidal signal is removed from the data by means of a 40-hour low-pass filter. Ref. [15] provides a more extensive description of this collected data set.

For the experiments presented in this paper, ADCP velocity measurements were extracted from the above data set for the month of June 2004. However, since the measurements have a very high resolution in the vertical and temporal dimensions, it is unrealistic to assimilate all of the velocity measurements, even with a very short cycle. For the entire time period of this experiment, there are roughly 1.75 million individual measurements (the u- and v- components of velocity are considered as 2 distinct measurements). Therefore, the data must be sampled at a prescribed temporal and vertical frequency. After careful examination of the entire dataset it was apparent that the dominant feature in time was inertial oscillations with a frequency of about 12 hours. To resolve these features, an assimilation frequency of 3 hrs was selected. In the vertical, it was apparent that on average the velocity profiles can be generalized with five or less layers. Therefore, at every 3 hr increment the two components of velocity are assimilated at up to five different depths at each of the 14 mooring locations. Therefore, the average number of measurements for a 12-hour cycle, for example, is 538. The five measurement depth locations are automatically selected independently at each assimilation time increment and at each horizontal mooring location. For each of these assimilated velocity profiles, the four strongest vertical gradients in velocity are determined and used to define the interfaces between the 5 layers. Then the measurement closest to the center of each layer is the one that is assimilated. If there is not a unique measurement in between two interfaces then that particular layer is ignored.

For each of the assimilated measurements, a measurement functional is created to translate the measurement from data space into the state space. The measurement functional can be as simple as a 1.0 at the grid point closest to the measurement location and zeros for the rest of the state. Since not all of the data are used, however, it is desired that the measurement functional represents the region of state space that encompasses the measurement and includes the area of the neglected data. Therefore, each measurement functional will include representation ± 1.5 hours of the measurement time, the entire layer that the measurement represents, and a 3 grid point horizontal radius surrounding the measurement. Each of these grid points, however, does not receive equal representation. A Gaussian function is used to distribute the influence in all 4-dimensions with the grid point closest to the measurement receiving the largest representation. Finally, all of the contributions to the measurement functional are scaled so that they sum up to one. In addition to the actual measurement and its associated measurement functional, a measurement error is required. Each measurement error is estimated to be the representation error that the assimilated measurement has with respect to the neglected data that the measurement is representing. This error is calculated by taking the RMS of the difference between the measurement value and each of the neglected data that falls within the ± 1.5 hour time window and the layer that the measurement represents.

D. The Cycling Representer Method

The representer method can be broken down into seven fundamental components: the TLM of the NL model, the adjoint of the TLM, the background, the measurements, the measurement functionals, the measurement errors, and of course the error covariances. An error covariance ought to be specified for each weak constraint variable that is believed to be in error. However, since error covariances are one of the more difficult quantities to accurately estimate and can be computationally expensive to include, it is not feasible to treat all variables in error as weak constraint; and for the variables that are treated as weak constraint, simplifications must be made in order to construct their error covariances. The assimilation results presented in this paper were computed by treating the initial conditions, the boundary conditions, and the interior solution of both velocity components as weak constraint. The covariance for each of these weak constraint variables was constructed with a constant variance in space and time, and correlations described by a Gaussian function in space and a moving average in time. An error of $5\,\mathrm{cm/s}$ was used for the initial and boundary conditions, and an error of $1\,\mathrm{cm/s}$ was used for the interior solution. The Gaussian functions used to estimate the correlation were formulated with an e-folding scale of 15 grid points in each of the 3 spatial dimensions and 2 hours in time.

The formulation of the cycling representer method that is used in this paper is the same as that described in great detail in [12], therefore it will not be prudent to repeat it here. Other than the seven fundamental components of the representer method mentioned above, the only difference between this assimilation system and the one used in [12] is how the background for each new cycle is computed. The background for each new cycle (other than the first) is computed by propagating the NL model forward using the final assimilated solution from the previous cycle as the initial condition and boundary conditions from the global solu-

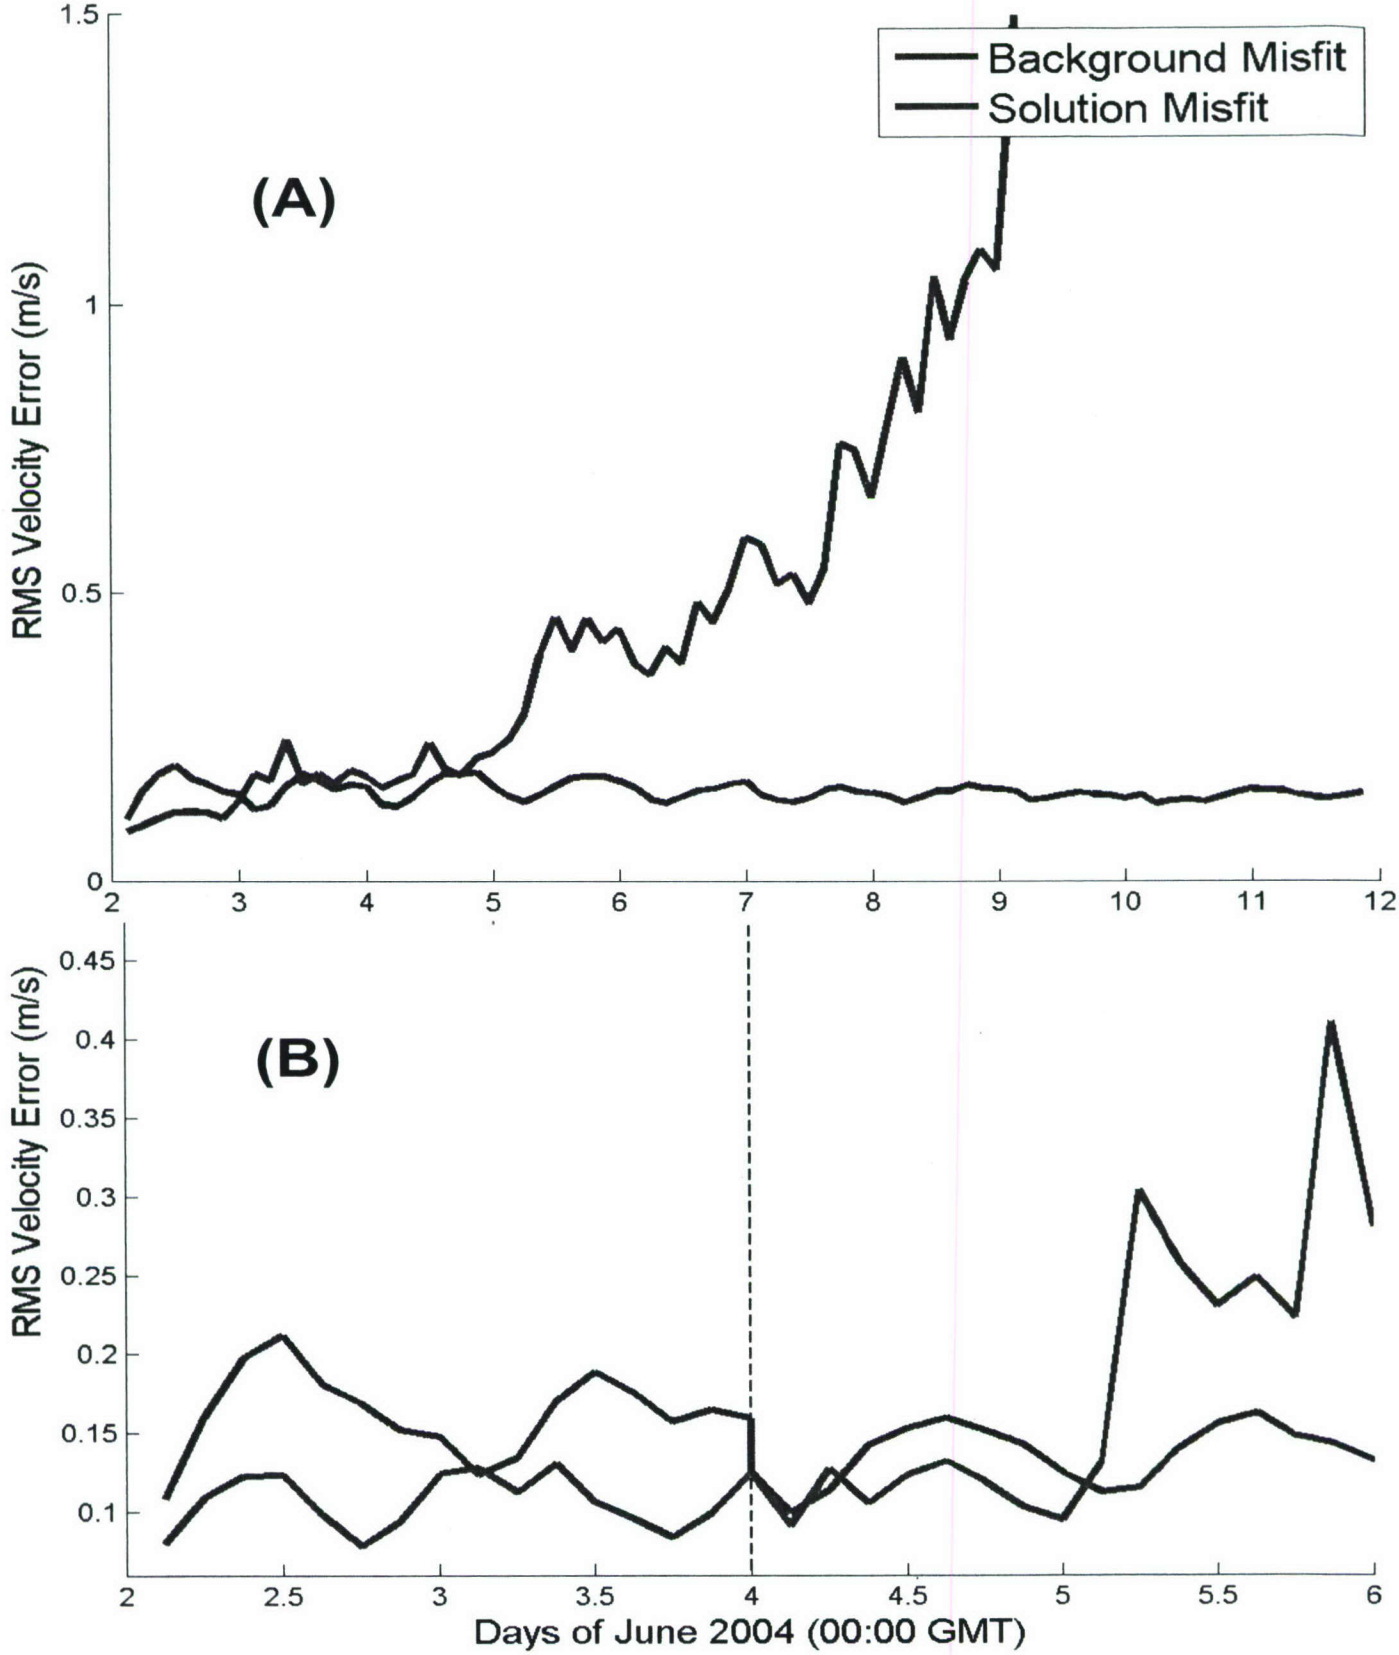

Fig. 3 $\mathrm{RMS}\left(L[\mathbf{u}]-\mathbf{d}\right)$ **is computed and plotted here at each 3-hour time stamp that data is assimilated.** L **is the measurement functional that translates the velocity field from state space to data space, and u represents velocity from the nonlinear background (blue) and the assimilated solution (red). (A) 10-day assimilation experiment. (B) Two 2-day cycle experiments, where the dashed line represents the break between cycles.**

tion. Due to the low-resolution of the global solution, there is a dynamic instability between the initial and boundary conditions. Therefore, in order for the new background to be stable its boundary conditions must have a smooth transition between the previous solution and global boundary conditions. To accomplish this transition, each assimilation is run an extra 3 hours over into the next cycle. Then a time weighted average is used on this solution and the global boundary conditions for these 3 hours to create smooth boundary conditions for the next background.

As is the case in [12], the representer method for each cycle is solved indirectly using an iterative conjugate gradient method (CG). This differs from the direct approach in that instead of computing representers for every measurement and inverting the representer matrix to obtain the representer coefficients, the indirect method uses the CG to iterate through the search directions in data space to converge upon the representer coefficients that minimize the cost function. The CG is considered converged when the norm of the residuals relative to the initial norm is less than 10^{-3}. For a well-conditioned system, the number of CG iterations is significantly less (typically about 10%) than the total number of measurements, therefore resulting in a substantial savings in computational cost compared to the direct method.

III. Results

A. Cycling Experiments

As a precursor to the cycling experiments, a long 10-day assimilation experiment is performed, and the resulting solution misfit (red) is plotted in Fig. 3a. The background misfit (blue) is also plotted for comparison. These misfits are computed at each of the 3-hour assimilation time increments and are the RMS of the difference between the data and the solution(background) acted upon by the same measurement functional described in Section 2c. The results in Fig. 3a reveal that there is a fairly consistent correlation between the accuracy of the assimilated solution and the TLM stability (Fig. 2). After the first day the assimilated solution begins to lose accuracy and by day 3 the errors in the solution begin to increase exponentially. Since the TLM is only stable for

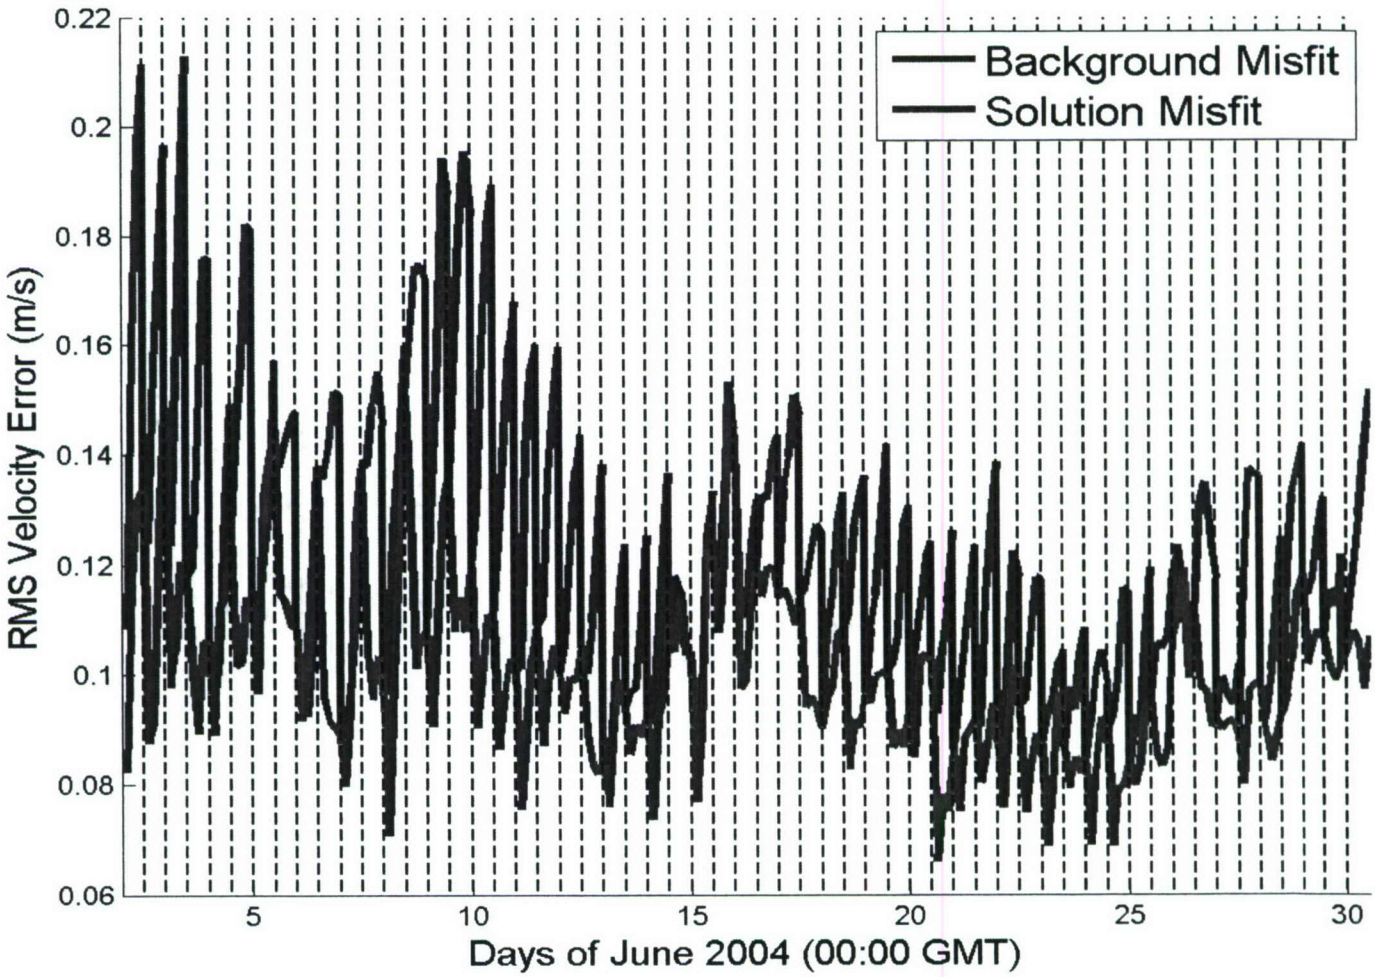

Fig. 4 This plot is similar to Fig. 3, except that 56 12-hour assimilation cycles are displayed.

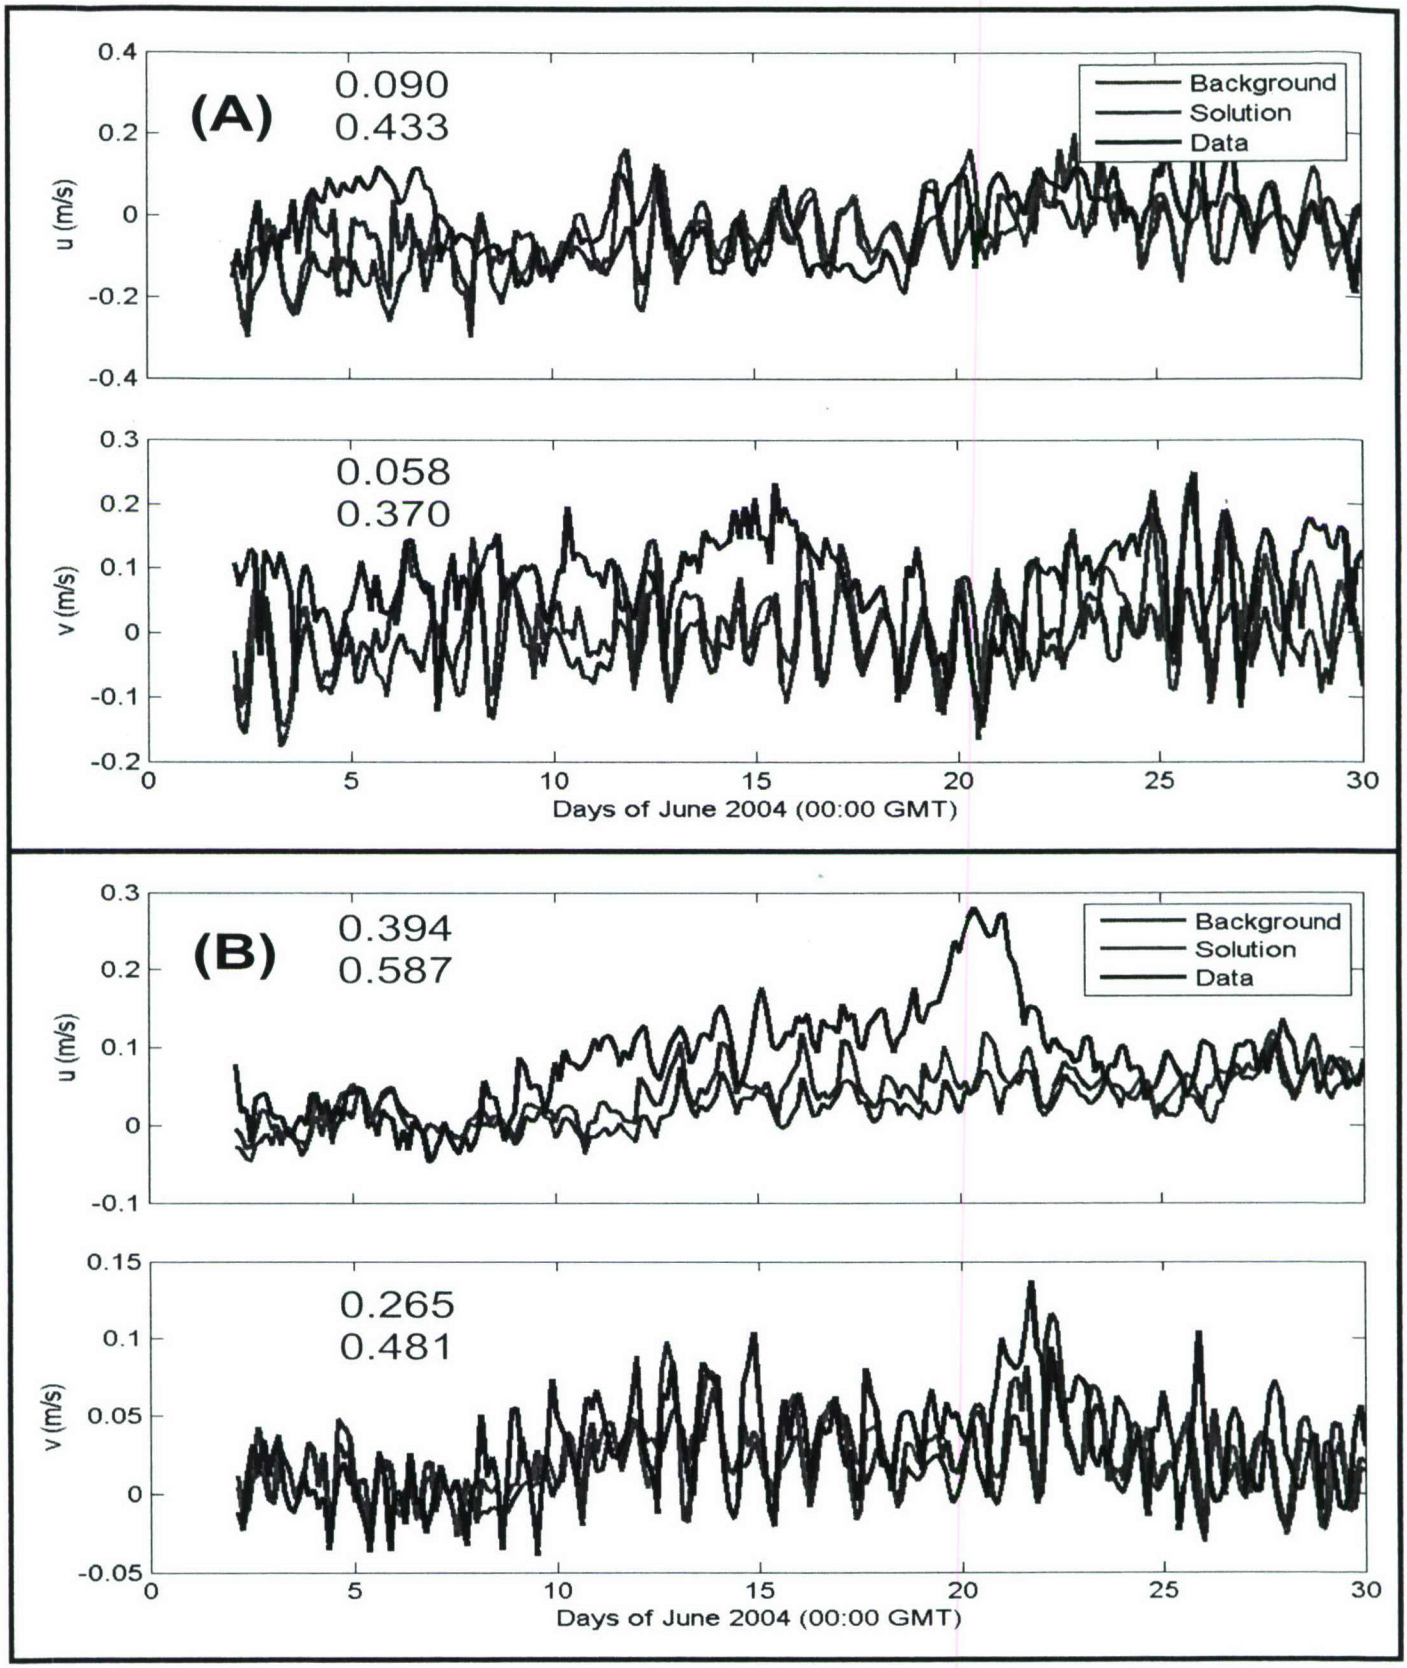

Fig. 5 Entire time series of velocity components of the data (black), solution (red), and background (blue) for the 12-hour cycle experiment at (A) the first layer of mooring location 7 and (B) the fifth layer of mooring location 13 (mooring locations are labeled in Fig. 1). The colored values represent the corresponding correlation coefficient between the background and data (blue), and the solution and data (red).

about a day, the assimilation cycle should be equal or shorter to this in order to ensure a stable and accurate solution. It should be noted that the CG in this assimilation experiment did not converge and was not going to converge; the problem was too ill-conditioned. The CG was stopped after 104 iterations of the conjugate gradient, and a final sweep was performed with the best estimate of representer coefficients that had been obtained up to that point.

Several different cycling experiments were performed to determine the optimal cycle time period that produces the solution with the overall best fit to the data. The optimal cycle time period needs to be short enough to ensure stability of the solution, and as long as possible to maximize the temporal correlation of the data and model. The first cycling experiment was performed using 2-day cycles. The misfit results are displayed in Fig. 3b and reveal that the first cycle did quite well, but the second cycle began to severely lose skill midway through the cycle. At the end of the second cycle the solution was too poor to provide a sufficient initial condition for the next background forecast (the forecast grew numerically unstable). The dashed black line in this figure represents the break in between cycles and the vertical blue line segment along this dashed line is a result of the background being reinitialized to the assimilated solution. It appears that a 2-day cycle time period is a little too long to ensure solution accuracy. This falls in line with the time frame of TLM stability.

Fig. 4 displays the resulting misfits for the next cycling experiment consisting of 12-hour cycles performed over 28 days. Overall, this assimilation experiment did quite well and there were no problems with providing a dynamically consistent solution as an initial condition for the forecasts. There are some ups and downs in the general trend of the solution misfit, but overall the misfit is improving through the cycles as will be verified in the next sub-section below. This plot is somewhat misleading in that it may initially appear that the solution misfit is not much better than the background misfit and that the background misfit is following the same general downward trend. One must remember that the background is continually being reinitialized to the solution at the end of each cycle, and since the time period is so long and there are so many cycles, there is not much space for the background misfit to grow. Despite this fact, the overall correlation coefficient between the solution and the data is 0.411, whereas the correlation coefficient between the background and data is 0.285. This improvement is very promising, especially when you factor in how difficult it is to correct a 4-dimensional velocity field towards a large dynamic dataset while maintaining dynamic stability of the model. To demonstrate this dynamic stability and the general improvement of the assimilated solution, time series of the velocity components for the background (blue), solution (red), and the data (black) are plotted in Fig. 5 for two selected locations: (a) the first layer of mooring 7 and (b) the fifth layer of mooring 13 (mooring locations are labeled in Fig. 1). The assimilated and background solutions are taken at the grid point and depth location that is closest to each velocity measurement. The colored values are the correlation coefficients between the background(solution) and the data for their corresponding time series.

The final cycling experiment consists of 12 1-day assimilation cycles and the misfits for this experiment are displayed in Fig. 6b. For comparison, the first 12 days of the 12-hour assimilation cycle experiment is plotted in Fig. 6a. After careful examination of these two plots, it is apparent that the 1-day cycle experiment is outperforming the 12-hour cycle experiment. The solution misfits in the 1-day experiment obtain lower values relative to the 12-hour cycle experiment, and the overall slope has a steeper downward trend. Also, in the 1-day cycle experiment there is a significant improvement in the background misfit. This is because there is a steep downward trend starting at the middle of each cycle. It is believed that this drastic change in the background misfit is due to the inertial oscillations, which are relatively strong in this region and have a frequency of about 12 hours. It appears that the longer 1-day cycles are able to better resolve the inertial oscillations and therefore produce a more accurate solution that better matches the observed flow field. This result demonstrates the importance of choosing a cycle time period that is long enough to include the important dynamic features that are prevalent in the region.

B. Performance of the Cycling Representer Method

In order to help analyze the performance of the Cycling Representer Method and gauge how well it might do if the system were allowed to continue to cycle indefinitely, several post-processing results are presented here for the entire 12-hour assimilation cycle experiment. To demonstrate the value of solving the Representer Method indirectly, Fig. 7 shows the evolution of the CG convergence for each of the 56 cycles (black lines). Overlaid on this wide spread of results is the best-fit exponential curve (thick red line). This curve conveys that the average total number of CG iterations required to reach convergence is about 62, which is roughly 12% of the average number of measurements in each cycle (538). This result corresponds to a substantial savings in computational cost and reveals that the assimilation problem is well-conditioned.

Fig. 8 displays the total number of required CG iterations (blue line) and the total cost of the cost function (green line) relative to the 12-hour cycles. Both of these results are quite noisy, but over the entire 28 time period their corresponding linear best-fits have a downward trend. The slight overall decrease in the number of required CG iterations as the system is cycled suggests that the assimilation problem for each consecutive cycle is becoming better conditioned, and the steep downward trend of the total cost reveals that the overall fit between the assimilated solution and the data, dynamics, initial conditions, and boundary conditions is improving as the cycling progresses forward. It is uncertain how long these trends would continue in this fashion if the cycling were to continue beyond the 28-day period of this experiment. This result, however, is promising in that this time period

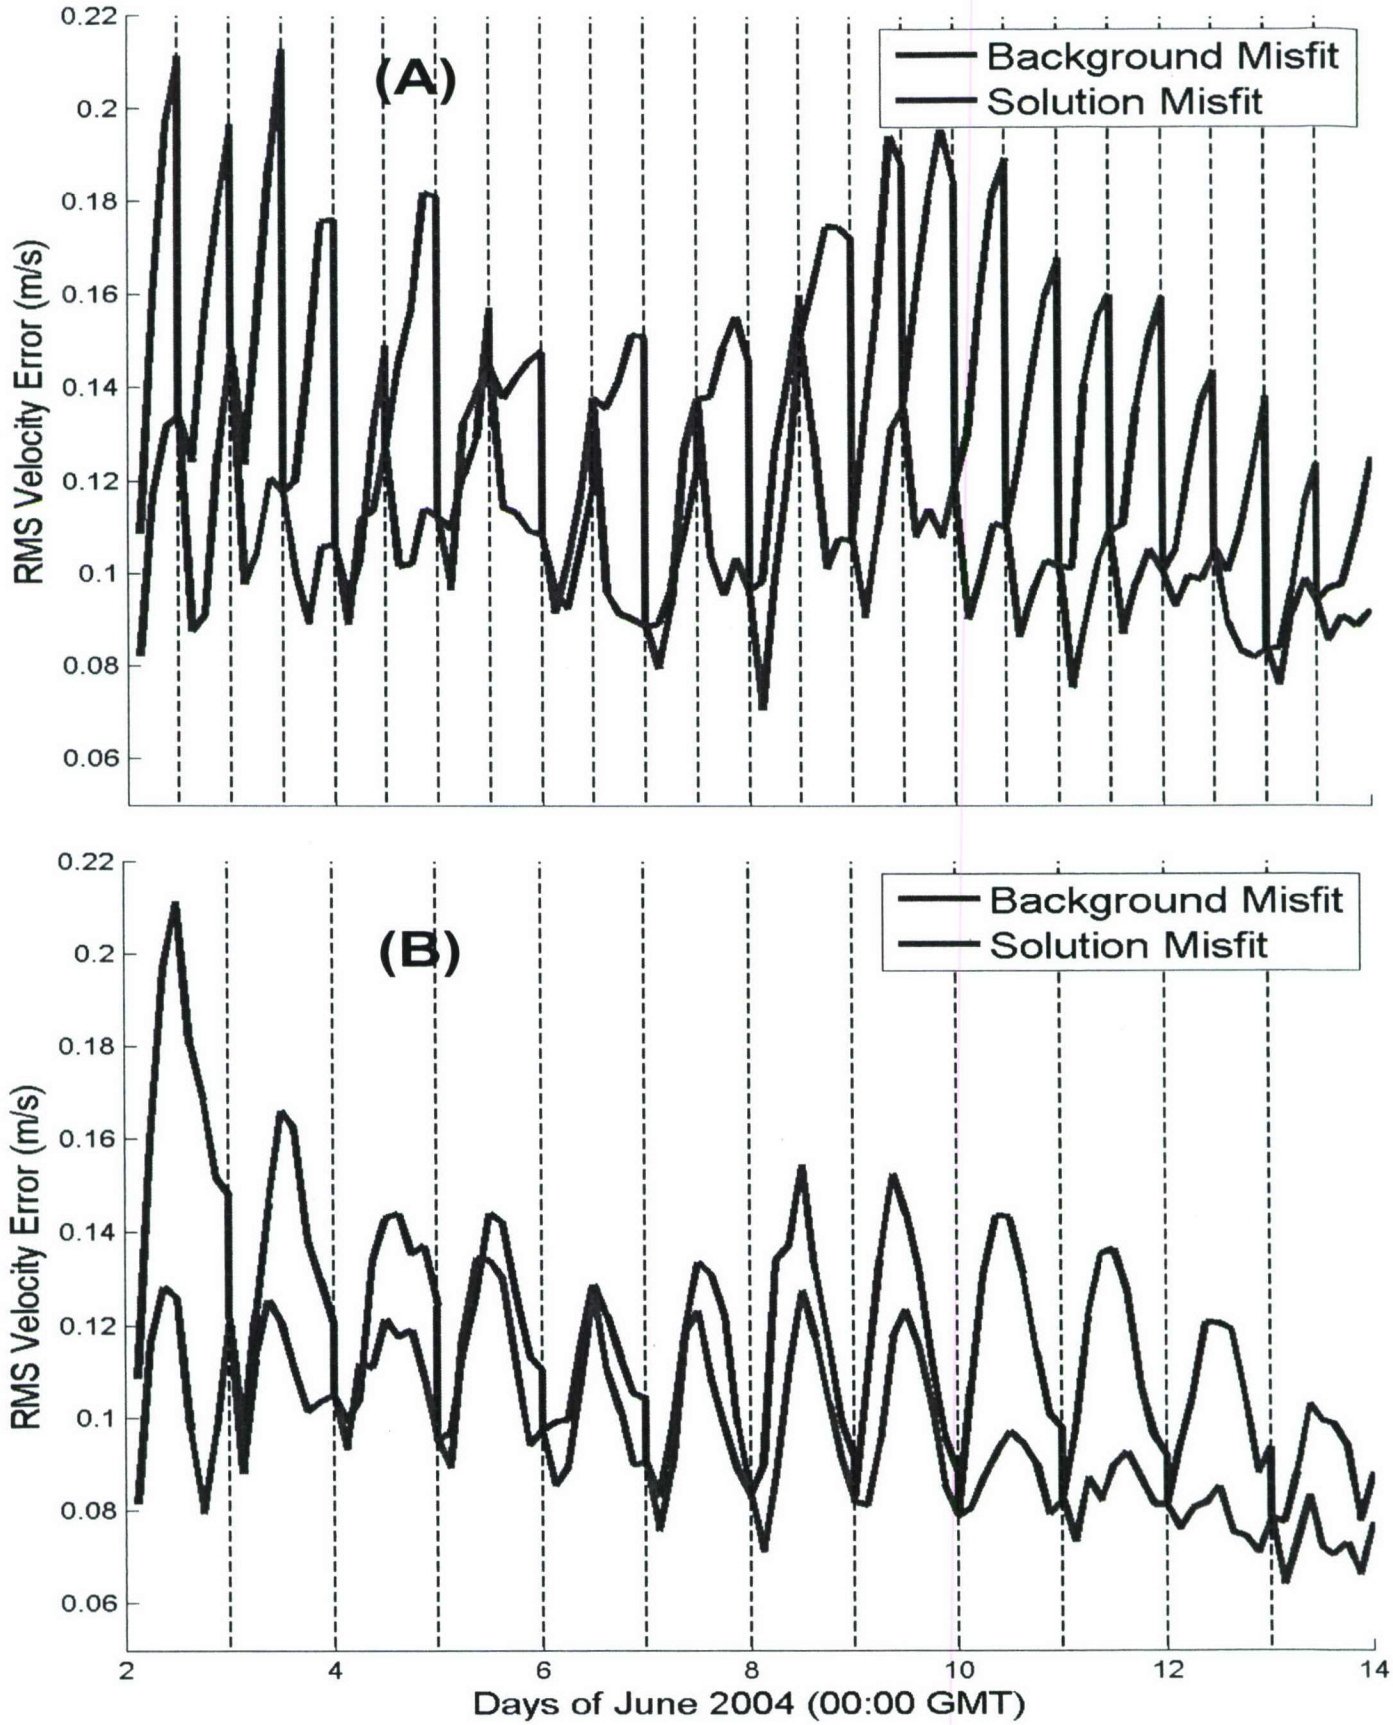

Fig. 6 (A) displays the first 12 days of the 12-hour cycle experiment from Fig. 4 (plotted for comparison) and (B) displays 12 1-day assimilation cycles.

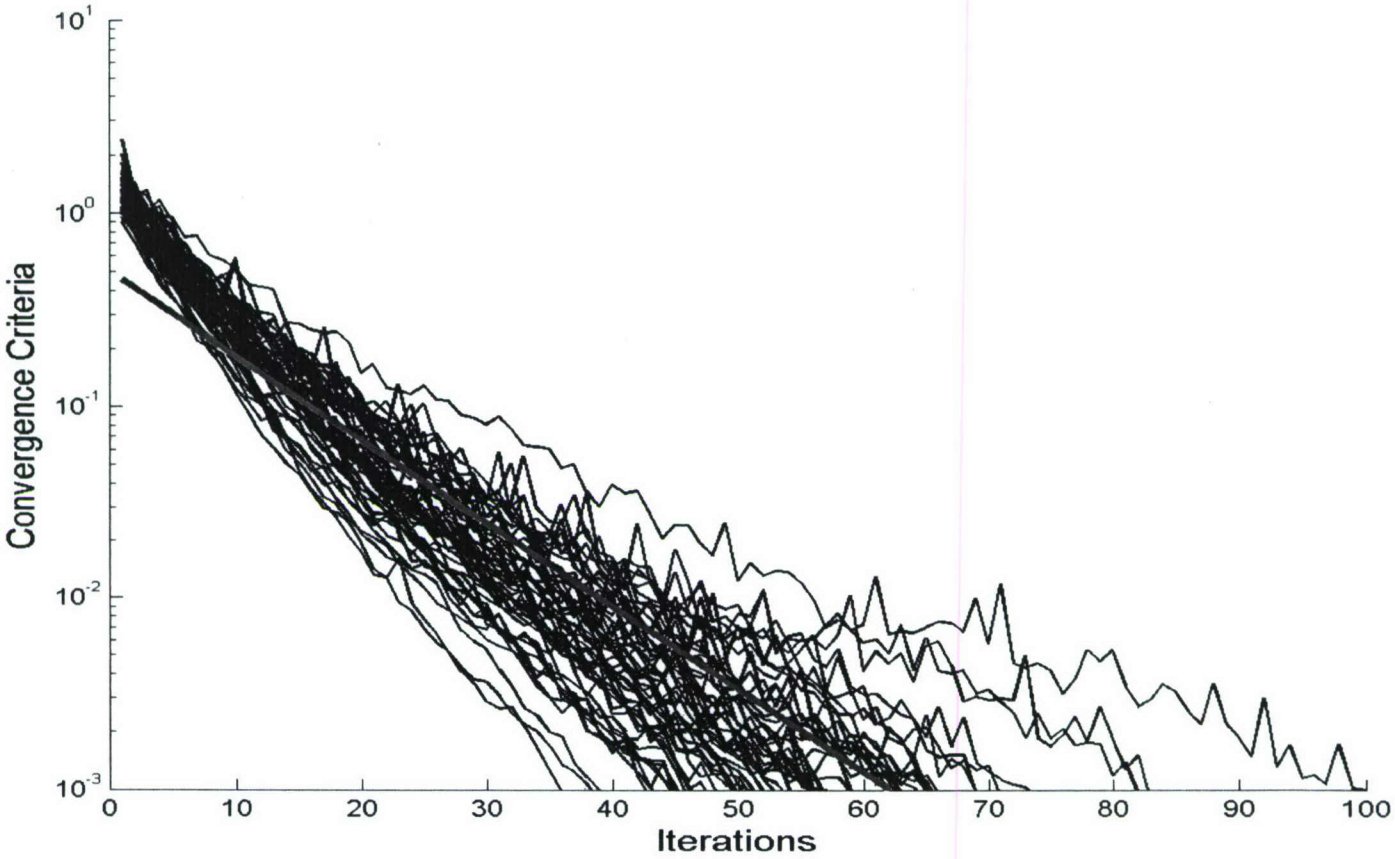

Fig. 7 The evolution of CG convergence criteria for each of the 56 12-hour assimilation cycles (black lines). The convergence criterion is calculated as the norm of the residuals relative to the initial norm and is assumed converged when it is less than or equal to 10^{-3}. The red line is the exponential best fit of the spread of CG convergence plots, showing that on average the 12-hour cycle converged in about 62 CG iterations. This is about 12% of the average number of measurements (538) in each cycle.

is long enough to contain a significant portion of the dynamic features that are prevalent in this region, and there is no reason to assume a drastic change in these trends.

IV. Conclusions

Even though the TLM of NCOM is only accurate for about a day, the cycling representer method can be used in conjunction with the CG to achieve an improved analysis of the flow field for what appears to be an indefinite period of time. Also, because the forecasts are reinitialized to the analysis field at the beginning of each cycle, both the forecast and analysis errors decrease as the system progresses through the cycles. Care needs to be taken though in selecting the appropriate time-frame of the cycle. If the cycle is too long and is beyond the range of TLM accuracy, the solution will end up diverging from the data and obstructing the continuity of the cycling system, as is the case in the 2-day cycle experiment. On the other hand, if the cycle is to short, valuable information can be lost due to the lack of temporal representation of important dynamic features thus causing inaccuracies in the analysis field, as is the case in the 12-hour cycle experiment. For the system described in this paper, the 1-day cycle experiment performed the best, and this cycle time-period seems to correspond with the outer accuracy limits of the TLM.

Overall, this paper demonstrates that the Cycling Representer Method can potentially be a valuable assimilation tool within an operational analysis/forecast system for coastal applications. One of the difficulties that would have to be overcome in order to achieve this operational status (other than improving the error covariances) is that the 1-day cycle period is only optimal for the specific application described in this paper, and most likely would not be the optimal choice if this assimilation system were to be applied to a different application. For example, if the assimilation system were applied to a different coastal region, or used a different background, or had a different resolution, etc..., then the optimal cycle time-period would have to be determined again. A possible solution to this problem could be to set up an autonomous system that is similar to what is used in [12]. i.e., a system that automatically defines the maximum time-frame of TLM accuracy as the amount of propagation time required before the difference between the TLM and the background surpass one standard deviation of the background. Then based on the results of this paper, the cycle time period can be set to this TLM accuracy criterion.

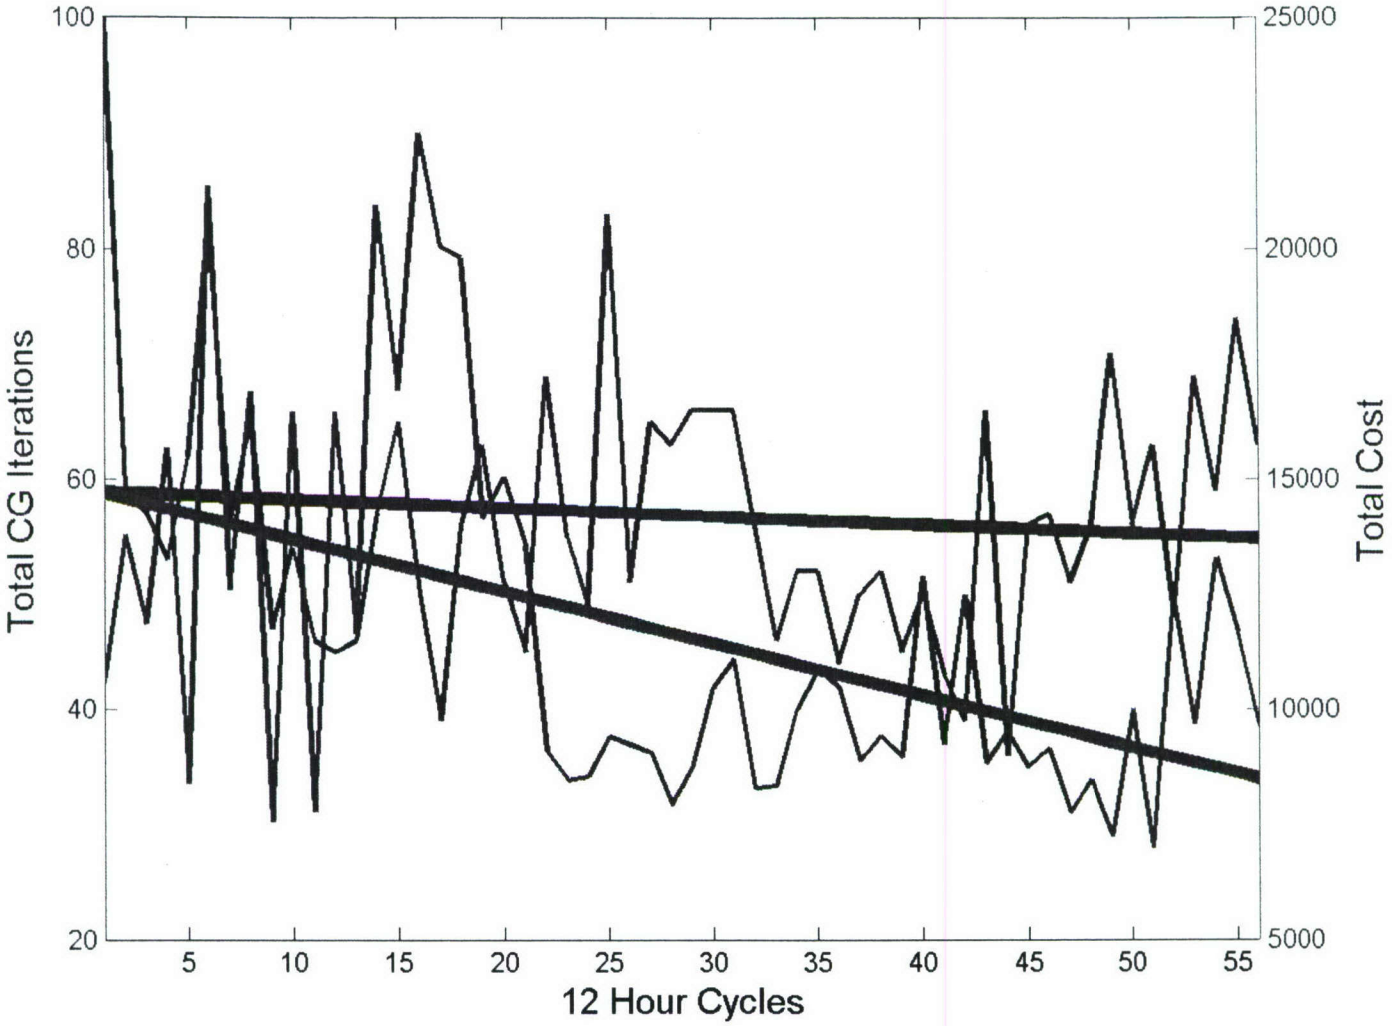

Fig. 8 Progression of the total number of required CG iterations (blue) and the total cost of the cost function (green) through the 12-hour cycles. Even though both of these results are quite noisy, their corresponding linear best fits (thick lines) have a downward trend as the number of cycles is increased.

ACKNOWLEDGMENT

This work is sponsored by the Office of Naval Research (Program Element number 0601153N) as part of the NRL project 'Shelf to Slope Energetics and Exchange Dynamics'.

REFERENCES

[1] C.N. Barron, A.B. Kara, P.J. Martin, R.C. Rhodes, L.F. Smedstad, "Formulation, implementation and examination of vertical coordinate choices in the Global Navy Coastal Ocean Model (NCOM)," *Ocean Modeling*, vol. 11, pp. 347-375, 2006.

[2] A.F. Bennett, B.S. Chua, and L.M. Leslie, "Generalized inversion of a global numerical weather prediction model," *Meteorol. Atmos. Phys.*, vol. 60, pp. 165-178, March 1996.

[3] A.F. Bennett, B.S. Chua, D. Harrison, and M.J. McPhaden, "Generalized inversion of tropical atmosphere-ocean data and a coupled model of the tropical Pacific. Part II: The 1995-96 La Nina and 1997-1998 El Nino," *Journal of Climate*, vol. 13, pp. 2770-2785, August 2000.

[4] A.F.Bennett, *Inverse Modeling of the Ocean and Atmosphere*. Cambridge University Press: Cambridge, 2002.

[5] A.F. Bennett, B.S. Chua, H.E. Ngodock, D.E. Harrison, and M.J. McPhaden, "Generalized inversion of the Gent-Cane model of the tropical Pacific with Tropical Atmosphere-Ocean (TAO) data," *Journal of Marine Research*, vol. 64, pp. 1-42, January 2006.

[6] T.F. Hogan and L.R. Brody, "Sensitivity studies of the Navy's global forecast model parameterizations and evaluation of improvements to NOGAPS," *Mon. Wea. Rev.*, vol. 121, pp. 2373-2395, August 1993.

[7] A.B. Kara, C.N. Barron, P.J. Martin, L.F. Smedstad, R.C. Rhodes, "Validation of interannual simulations from the 1/8° global Navy Coastal Ocean Model (NCOM)," *Ocean Modeling*, vol. 11, pp. 376-398, 2006.

[8] A.M. Moore, H.G. Arango, E.D. Lorenzo, B.D. Cornuelle, A.J. Miller, and D.J. Neilson, "A comprehensive ocean prediction and analysis system based on the tangent linear and adjoint of a regional ocean model," *Ocean Modelling*, vol. 7, pp. 227-258, 2004.

[9] J.C. Muccino and H. Luo, "Picard iterations for a finite element shallow water equation model," *Ocean Modelling*, vol. 10, pp. 316-341, 2005.

[10] H.E. Ngodock, B.S. Chua, and A.F. Bennett, "Generalized inverse of a reduced gravity primitive equation ocean model and tropical atmosphere-ocean data," *Mon. Wea. Rev.*, 128, pp. 1757-1777, June 2000.

[11] H.E. Ngodock, G.A. Jacobs, and M. Chen, "The representer method, the ensemble Kalman filter and the ensemble Kalman smoother: a comparison study using a nonlinear reduced gravity ocean model," *Ocean Modelling*, vol. 12, pp. 378-400, 2006.
[12] H.E. Ngodock, S.R. Smith, and G.A. Jacobs, "Cycling the representer algorithm for variational data assimilation with the Lorenz attractor," *Mon. Wea. Rev.*, vol. 135, pp. 373-386, February 2007.
[13] H.E. Ngodock, S.R. Smith, and G.A. Jacobs, "Cycling the representer algorithm for variational data assimilation with a nonlinear reduced gravity ocean model," *Ocean Modeling*, in press.
[14] S.R. Smith and G.A Jacobs, "Seasonal circulation fields in the northern Gulf of Mexico calculated by assimilating current meter, shipboard ADCP, and drifter data simultaneously with the shallow water equations," *Continental Shelf Research*, vol. 25, pp.157-183, 2005.
[15] W.J. Teague, E. Jarosz, M.R. Carnes, D.A. Mitchell, and P.J. Hogan, "Low-frequency current variability observed at the shelfbreak in the northeastern Gulf of Mexico: May-October, 2004," *Continental Shelf Research*, vol. 26, pp.2559-2582, 2006.
[16] L. Xu and R. Daley, "Towards a true 4-dimensional data assimilation algorithm: application of a cycling representer algorithm to a simple transport problem," *Tellus*, vol. 52A, pp. 109-128, March 2000.
[17] L. Xu and R. Daley, "Data assimilation with a barotropically unstable shallow water system using representer algorithms," *Tellus*, vol. 54A, pp. 125-137, March 2002.